

# Approximate Equilibria in Games with Few Players\*

Patrick Briest  
 Dept. of Computer Science  
 University of Liverpool, U.K.  
 Patrick.Briest@liverpool.ac.uk

Paul W. Goldberg  
 Dept. of Computer Science  
 University of Liverpool, U.K.  
 P.W.Goldberg@liverpool.ac.uk

Heiko Röglin  
 Dept. of Computer Science  
 RWTH Aachen, Germany  
 roeglin@cs.rwth-aachen.de

April 29, 2008

## Abstract

We study the problem of computing approximate Nash equilibria ( $\epsilon$ -Nash equilibria) in normal form games, where the number of players is a small constant. We consider the approach of looking for solutions with constant support size. It is known from recent work that in the 2-player case, a  $\frac{1}{2}$ -Nash equilibrium can be easily found, but in general one cannot achieve a smaller value of  $\epsilon$  than  $\frac{1}{2}$ . In this paper we extend those results to the  $k$ -player case, and find that  $\epsilon = 1 - \frac{1}{k}$  is feasible, but cannot be improved upon. We show how stronger results for the 2-player case may be used in order to slightly improve upon the  $\epsilon = 1 - \frac{1}{k}$  obtained in the  $k$ -player case.

## 1 Introduction

A game in normal form has  $k$  players, and for each player  $p$  a set  $S^p$  of pure strategies. In this paper we assume that all sets  $S^p$  are of the same size  $n$ . The set  $S$  of pure strategy profiles is the cartesian product of the sets  $S^p$ . For each player  $p$  and each  $s \in S$ , the game has an associated value  $u_s^p$ , being the utility or payoff to player  $p$  if all players choose  $s$ . Note that the number of quantities needed to specify the game is  $k \cdot n^k$ , so we must take  $k$  to be a constant, for the game's description to be of size polynomial in  $n$ .

A mixed strategy for player  $p$  is a probability distribution over  $S^p$ . Suppose that each player  $p$  chooses distribution  $D^p$  over  $S^p$ . Let us define a player's *regret* to be the highest payoff he could obtain by choosing a best response to the other players' mixed strategies, minus his actual expected payoff. Then, a *Nash equilibrium* is a set of  $D^p$ 's for which all players' regrets are zero. We say that the distributions  $D^p$  form an  $\epsilon$ -*Nash equilibrium* provided that all players' regrets are at most  $\epsilon$ .

---

\*This work was supported by DFG grant VO 889/2, EPSRC Grant GR/T07343/02, and by the EU within the 6th Framework Programme under contract 001907 (DELIS).

## 1.1 Recent work

Due to the apparent difficulty of computing Nash equilibria exactly [4, 5, 6], recent work has addressed the question of polynomial-time computability of approximate Nash equilibria, in particular the  $\epsilon$ -Nash equilibria defined above. The effort has mostly addressed 2-player games, and the general question is, for which values of  $\epsilon$  can  $\epsilon$ -Nash equilibria be found in polynomial time? The payoffs  $u_s^p$  in the games are restricted to lie in the range  $[0, 1]$ , since otherwise the payoffs (and associated values of  $\epsilon$ ) could be rescaled arbitrarily. Recent papers include [2, 13], which show how to efficiently compute 2-player  $\epsilon$ -Nash equilibria for  $\epsilon = 0.36392$  and  $0.3393$  respectively. We do not know of similar work so far that addresses the  $k$ -player case considered here.

The *support* of a probability distribution is the number of elements of its domain that have non-zero probability. Solutions of games having small support (that is, players' mixed strategies do not give positive probability to many of the pure strategies) are attractive for two reasons. From the perspective of modelling a plausible outcome, we expect a participant in a game to prefer a simple behaviour. Also, if constant-support solutions exist, then if the number of players is constant, we can find them in polynomial time by brute force. For example, this approach is used by Bárány et al. [3] to find Nash equilibria in random games. Our results show that additional players makes it genuinely more intractable to find a satisfactory solution; if we restrict the support sizes to any constant then the worst-case regret of some player increases.

It is known that for 2 players, if we restrict ourselves to solutions with constant support, then there is a lower bound of  $\frac{1}{2}$  on the best  $\epsilon$  that can be achieved [10], and this lower bound is achieved by a very simple algorithm of Daskalakis, Mehta and Papadimitriou [7].

Regarding the issue of how large the support size needs to be in order to allow an  $\epsilon$ -Nash equilibrium, it is shown by Althöfer [1] and Lipton et al. [12] that  $\log(n)/\epsilon^2$  support size is sufficient, and [12] point out that the general method works for any constant number  $k$  of players. It follows that the brute-force “support enumeration” algorithm takes time  $O(n^{\log n})$  for any  $\epsilon > 0$ .

Very recently, Hémon et al. [11] studied independently from our work the problem of computing approximate Nash equilibria in games with more than two players. The results they obtain are very similar to the results presented in this paper. In particular, they also show that for games with  $k$  players a  $(1 - \frac{1}{k})$ -Nash equilibrium with constant support size can be computed efficiently and that this is the best possible which can be achieved with constant support strategies. Additionally, they also show that support size  $O(\log(n)/\epsilon^2)$  is sufficient for computing  $\epsilon$ -approximate Nash equilibria in the additive and multiplicative sense.

## 1.2 Our Results

We give a simple algorithm, essentially an extension of [7] from the 2-player case to the  $k$ -player case, that finds  $(1 - \frac{1}{k})$ -Nash equilibria in which each player's mixed strategy has support size at most 2. Notice that this result becomes quite weak as  $k$  increases, given that any set of mixed strategies constitutes a 1-Nash equilibrium (recall that we assume all payoffs lie in the range  $[0, 1]$ ). However, we also show that one can do no

better for constant support strategies. The argument is a kind of generalisation of the one of [10] that gives the result for the 2-player case.

For unrestricted strategies, we give a method that uses any algorithm for the 2-player case as a component, and provides slightly better values of  $\epsilon$  than the above. Finally, we show how the bounded differences inequality can be used to give a simplified proof of a result of [12], that for a constant number of players,  $O(\log n)$  support size is sufficient for finding approximate equilibria. This allows us to deduce that if the number of players is less than  $\sqrt{n}$ , there is a subexponential algorithm for finding approximate Nash equilibria.

## 2 Details of Results

### 2.1 Solutions with Constant Support

In this section we generalise the results of [7, 10] from  $\epsilon = \frac{1}{2}$  in the 2-player case, to  $\epsilon = 1 - \frac{1}{k}$  in the  $k$ -player case.

**Definition 1.** *Define a winner-takes-all game to be one in which, for any combination of pure strategies, one player obtains a payoff of 1, and the other players obtain payoffs of 0.*

In the 2-player case, winner-takes-all games are win-lose games where the payoffs sum to 1. Our generalisation of the lower bound of [10] follows their approach in that we generate a random winner-takes-all game and show that for support sizes limited by some constant  $t$ , for any  $\epsilon < 1 - \frac{1}{k}$ , if the game is large enough, there will be positive probability that any solution with supports at most  $t$  is not  $\epsilon$ -Nash.

**Definition 2.** *Define the total support of a mixed-strategy profile to be the sum over all players  $p$ , of the number of strategies of  $p$  that  $p$  assigns non-zero probability.*

**Theorem 3.** *Let  $t$  and  $k$  be any positive integers, and let  $\epsilon = 1 - \frac{1}{k}$ . There exist games having  $k$  players such that in any  $\epsilon$ -Nash equilibrium, there must be at least one player whose mixed strategy has support greater than  $t$ . (In particular, there exist games with  $n$  strategies for each player such that in any  $\epsilon$ -NE at least one player must have a strategy with support  $\Omega(\sqrt[k-1]{\log n})$ .)*

*Proof.* We construct a  $k$ -player winner-takes-all game  $\mathcal{G}$  uniformly at random as follows. Let each player have  $n$  pure strategies (a suitable value of  $n$  will be identified later). For each combination of pure strategies, choose one of the  $k$  players uniformly at random and let that player have an associated payoff of 1, while the remaining players have payoffs of 0.

We prove that for large enough  $n$ , with positive probability, the resulting game has the desired property.

Consider strategy profiles with total support  $t$ . There are less than  $\binom{kn}{t}$  possible support sets.

We choose  $n$  large enough such that for any individual support set of total size  $t$ , the probability (with respect to random generation of  $\mathcal{G}$ ) that all players have regret less than  $\epsilon$  is less than  $1/\binom{kn}{t}$ , so that we can apply a union bound.

Given any mixed-strategy profile, we know that there exists at least one player who has an expected payoff of at most  $1/k$ , since the payoffs always sum to 1. Fix a support set  $S$  of total size  $t$ . We show that all players (including in particular the one with lowest expected payoff) have (with positive probability) a pure strategy that, if they unilaterally defected to it, would give them payoff 1 (for any mixed profile using  $S$ ).

For each player  $p$ , there are  $n$  pure strategies to choose from. In order for some given strategy  $s^p$  (available to player  $p$ ) to have expected payoff of 1 with respect to the other players' strategies, a sufficient condition is that the payoff entries corresponding to  $s^p$  and the  $\leq t - 1$  strategies in use by the other players, should let player  $p$  win. There are fewer than  $t^{k-1}$  pure-strategy combinations that the other players can select with non-zero probability.

The probability that  $s^p$  wins against a fixed pure-strategy combination is  $1/k$ , so the probability that  $s^p$  wins against all pure-strategy combinations of the other players is  $k^{-(t^{k-1})}$ . Hence, the probability that player  $p$  has no strategy that guarantees him a payoff of 1 is

$$\left(1 - k^{-(t^{k-1})}\right)^n.$$

The probability that there exists a player  $p$  without a winning strategy is at most

$$k \cdot \left(1 - k^{-(t^{k-1})}\right)^n.$$

The probability that there exists a support set for which there exists a player  $p$  without a winning strategy is at most

$$\binom{kn}{t} \cdot k \cdot \left(1 - k^{-(t^{k-1})}\right)^n.$$

We can upper bound this probability by

$$k \cdot (kn)^t \cdot \left(1 - k^{-(t^{k-1})}\right)^n.$$

Let us set  $t = \sqrt[k-1]{a \log_k(n)}$  for some constant  $a > 0$ . Then the term simplifies to

$$k \cdot (kn)^{\sqrt[k-1]{a \log_k(n)}} \cdot \left(1 - n^{-a}\right)^n. \quad (1)$$

For  $a = 1/2$ , we can estimate  $(1 - n^{-a})^n$  by  $e^{-\sqrt{n}}$  and hence (1) tends to zero when  $n$  tends to infinity.

This ensures that if the size of the support set is bounded from above by  $\sqrt[k-1]{\log_k(n)}/2$ , then, for large enough  $n$ , with positive probability each player has a pure strategy with payoff 1, regardless of the choice of the support set.

Hence, we require a total support size of  $\Omega(\sqrt[k-1]{\log n})$  to ensure an approximation performance of better than  $1 - 1/k$ , which for two players agrees with the upper bound of [12].  $\square$

**Corollary 4.** *If we restrict strategies to have constant support, there is a lower bound of  $1 - \frac{1}{k}$  on the approximation quality that we can guarantee.*

Note that the above is a generalisation of a lower bound of [1, 10] from the 2-player case to the  $k$ -player case. We show that this lower bound is optimal by giving an algorithm (which is an extension of an algorithm of [7]) that finds a  $1 - \frac{1}{k}$ -approximate Nash equilibrium where each player has support size at most 2.

**Theorem 5.** *Let  $\mathcal{G}$  be a  $k$ -player game with  $n$  strategies per player. Let  $\epsilon = 1 - \frac{1}{k}$ . There is a mixed strategy profile in which each player has support at most 2 which constitutes an  $\epsilon$ -Nash equilibrium.*

*Proof.* The following algorithm achieves the stated objective.

1. Number the players  $1, \dots, k$ . For  $i = 1, \dots, k - 1$  in increasing order, player  $i$  allocates a probability of  $1 - \frac{1}{k+1-i}$  to an arbitrary pure strategy  $s_i$ .
2. Player  $k$  allocates probability 1 to a pure best response  $b_k$  to the strategy combination  $(s_1, \dots, s_{k-1})$  of players  $1, \dots, k - 1$ .
3. For  $i = k - 1, \dots, 1$  in decreasing order, player  $i$  allocates his remaining probability  $\frac{1}{k+1-i}$  to a pure best response  $b_i$  to the mixed strategy formed so far, i.e., to the strategy  $(s_1, \dots, s_{i-1}, r_{i+1}, \dots, r_k)$ , where  $r_j$  denotes the mixed strategy  $(1 - \frac{1}{k+1-j}) \cdot s_j + (\frac{1}{k+1-j}) \cdot b_j$ .

In order to prove the theorem, we compute the regret of a player  $i \in \{1, \dots, k\}$ . To avoid case distinctions, we denote by  $s_k$  an arbitrary pure strategy of player  $k$ . If player  $i$  plays strategy  $s_i$ , which happens with probability  $1 - \frac{1}{k+1-i}$ , his regret can only be bounded by 1. If player  $i$  plays strategy  $b_i$ , then his regret is 0 if all players  $j < i$  play strategy  $s_j$  and his regret can be as bad as 1 otherwise. Altogether, this implies that the regret of player  $i$  is bounded by

$$\begin{aligned}
& \Pr[i \text{ plays } s_i] + \Pr[i \text{ plays } b_i] \cdot \Pr[\exists j < i : j \text{ plays } b_j] \\
&= \left(1 - \frac{1}{k+1-i}\right) + \frac{1}{k+1-i} \left(1 - \prod_{j=1}^{i-1} \left(1 - \frac{1}{k+1-j}\right)\right) \\
&= 1 - \frac{1}{k} . \quad \square
\end{aligned}$$

## 2.2 Solutions with non-constant support

The general issue of interest is the question of what approximation guarantee can we obtain for a polynomial-time algorithm for  $k$ -player approximate Nash equilibrium. To what extent can we improve on the above (weak) result when we allow unrestricted mixed strategies? We do not have a very substantial improvement, but the following shows how the trick of [7] can be combined with better 2-player algorithms to get an improvement over the  $(1 - \frac{1}{k})$ -result that the above would yield for  $k$  players.

**Theorem 6.** *If 2-player  $\epsilon$ -approximate Nash equilibria can be found in polynomial time, then  $k$ -player  $\frac{(k-2)-(k-3)\epsilon}{(k-1)-(k-2)\epsilon}$ -Nash equilibria can be found in polynomial time.*

As a corollary, the approximation guarantee of 0.3393 associated with the algorithm of [13] gives a 3-player algorithm with an approximation guarantee of 0.6022.

*Proof.* We prove the theorem by induction on the number of players. Let  $\delta_k = \frac{(k-2)-(k-3)\epsilon}{(k-1)-(k-2)\epsilon}$ . For  $k = 2$ , which is the induction basis,  $\delta_2 = \epsilon$ . Now assume that there is an algorithm  $\mathcal{A}_{k-1}$  with approximation guarantee  $\delta_{k-1}$  for  $(k-1)$ -player games. The following algorithm for  $k$ -player games achieves the stated approximation guarantee  $\delta_k$ .

1. Player 1 allocates a probability of  $1/(2 - \delta_{k-1})$  to some arbitrary pure strategy  $s$ .
2. Players  $2, \dots, k$  apply  $\mathcal{A}_{k-1}$  to the  $(k-1)$ -player game that results from letting player 1 play  $s$ .
3. Player 1 allocates his remaining probability to a pure best response  $b$  to the strategies of players  $2, \dots, k$  chosen in step 2.

A player  $i \in \{2, \dots, k\}$  has regret at most  $\delta_{k-1}$  if player 1 plays strategy  $s$  and his regret can be as bad as 1 otherwise. Hence, the regret of a player  $i \in \{2, \dots, k\}$  can be bounded by

$$\left(\frac{1}{2 - \delta_{k-1}}\right) \cdot \delta_{k-1} + \left(1 - \frac{1}{2 - \delta_{k-1}}\right) = \frac{1}{2 - \delta_{k-1}}.$$

Player 1 has no regret when playing  $b$  and his regret can be as bad as 1 when playing  $s$ . This implies that also the regret of player 1 can be bounded by  $1/(2 - \delta_{k-1})$ . A simple calculation shows

$$\frac{1}{2 - \delta_{k-1}} = \frac{1}{2 - \frac{(k-3)-(k-4)\epsilon}{(k-2)-(k-3)\epsilon}} = \frac{(k-2) - (k-3)\epsilon}{(k-1) - (k-2)\epsilon} = \delta_k,$$

as desired.  $\square$

Finally, we use the bounded differences inequality to give a simplified proof of a result of [12].

**Theorem 7.** *Let  $\mathcal{G}$  be a  $k$ -player game in which each player has  $n$  pure strategies. Let  $\epsilon > 0$ . Then there exists a mixed strategy profile where each player has support  $k^2 \log(kn)/2\epsilon^2$  that constitutes an  $\epsilon$ -Nash equilibrium. In particular, the players' distributions are empirical distributions over multisets of size  $k^2 \log(kn)/2\epsilon^2$ .*

*Proof.* The bounded differences inequality (see e.g. [9] p. 8) is the following. Let  $A$  be a set. Let  $g : A^N \rightarrow \mathbf{R}$  be a function with the *bounded difference* property: for all  $i \in \{1, \dots, n\}$  and for all  $x_1, \dots, x_N, x'_i \in A$

$$|g(x_1, \dots, x_N) - g(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_N)| \leq c_i.$$

Suppose  $X_1, \dots, X_N$  are independent random variables over  $A$ . The bounded difference inequality states that for all  $t > 0$ ,

$$\begin{aligned} & \Pr[g(X_1, \dots, X_N) - \mathbf{E}g(X_1, \dots, X_N) \geq t] \\ & \leq \exp\left(\frac{-2t^2}{\sum_{i=1}^N c_i^2}\right), \end{aligned}$$

and

$$\begin{aligned} & \Pr[\mathbf{E}g(X_1, \dots, X_N) - g(X_1, \dots, X_N) \geq t] \\ & \leq \exp\left(\frac{-2t^2}{\sum_{i=1}^N c_i^2}\right). \end{aligned}$$

We apply the above as follows. Let  $A = [n]^k$ , the set of all pure profiles of  $k$ -person games where each player has  $n$  strategies. The  $X_i$  are samples from some fixed Nash equilibrium. Let  $g_{i,j}$  be the payoff to player  $i$  for using strategy  $j$  subject to all players using the mixture of  $N$  strategies obtained by taking, for each player  $p$ , the  $p$ -th entries of each  $X_m$  ( $m \in [N]$ ), and using the uniform distribution on that multiset.

Note that for the  $g_{i,j}$  functions,  $c_i \leq k/N$ .

We want the right-hand sides to be less than  $1/2kn$ , so that by a union bound, for each  $i$  and  $j$ , we have the payoff for player  $i$  using strategy  $j$  being close to expected. Thus, we want  $N$  such that

$$\exp(-2\epsilon^2/(N(k^2/N^2))) \leq \frac{1}{2kn}.$$

Solving for  $N$  we get

$$N \geq \frac{k^2 \log(2kn)}{2\epsilon^2}. \quad \square$$

The theorem indicates that support enumeration is a subexponential algorithm provided that  $k = o(n^{1/2})$  (a brute-force search has complexity  $O(n^{kN})$  which is subexponential for  $k = O(n^{1/2-\zeta})$ .) Let us remark, that for  $k = o(n^{1/3})$  support enumeration is not only subexponential in the input size  $n^k$  but also in the number  $n$  of different strategies of the players. Hence, in this case, the running time is subexponential even for more succinctly represented games.

### 3 Conclusions

It seems that only very weak approximation performance is possible with constant-support strategies. Work on the special case of 2-player games has mostly addressed this problem by constructing linear programs whose solutions have useful properties (and are typically non-constant support strategies). For more than 2 players this approach no longer works; the corresponding constraints are no longer linear. The challenge seems to be to find alternative ways of describing mixed strategies (of more than constant support) that have desirable properties.

### References

- [1] I. Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199:339–355, 1994.
- [2] H. Bosse, J. Byrka and E. Markakis. New Algorithms for Approximate Nash Equilibria in Bimatrix Games. In *Proc. of the 3rd Int. Workshop on Internet and Network Economics (WINE)*, pages 17–29, 2007.

- [3] I. Bárány, S. Vempala and A. Vetta. Nash Equilibria in Random Games. In *Proc. of the 46th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 123–131, 2005.
- [4] X. Chen and X. Deng. Settling the Complexity of 2-Player Nash-Equilibrium. In *Proc. of the 47th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 261–272, 2006.
- [5] X. Chen, X. Deng and S. Teng. Computing Nash Equilibria: Approximation and Smoothed Complexity. In *Proc. of the 47th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 603–612, 2006.
- [6] C. Daskalakis, P. W. Goldberg and C. H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. In *Proc. of the 38th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 71–78, 2006.
- [7] K. Daskalakis, A. Mehta and C.H. Papadimitriou. A Note on Approximate Nash Equilibria. In *Proc. of the 2nd Int. Workshop on Internet and Network Economics (WINE)*, pages 297–306, 2006.
- [8] K. Daskalakis, A. Mehta and C.H. Papadimitriou. Progress in Approximate Nash Equilibria. In *Proc. of the 8th ACM Conference on Electronic Commerce (EC)*, pages 355–358, 2007.
- [9] L. Devroye and G. Lugosi. *Combinatorial Methods in Density Estimation*. Springer, 2001.
- [10] T. Feder, H. Nazerzadeh and A. Saberi. Approximating Nash Equilibria Using Small-support Strategies. In *Proc. of the 8th ACM Conference on Electronic Commerce (EC)*, pages 352–354, 2007.
- [11] S. Hémon, M. de Rougemont and M. Santha. Approximate Nash Equilibria for multi-players Games. In *Proc. of the 1st Int. Symp. on Algorithmic Game Theory*, 2008. to appear.
- [12] R. Lipton, V. Markakis and A. Mehta. Playing Large Games using Simple Strategies. In *Proc. of the 4th ACM Conference on Electronic Commerce (EC)*, pages 36–41, 2003.
- [13] H. Tsaknakis and P.G. Spirakis. An Optimization Approach for Approximate Nash Equilibria. In *Proc. of the 3rd Int. Workshop on Internet and Network Economics (WINE)*, pages 42–56, 2007.